

Stored Program Control No. 1A Store

**By W. A. BAKER, G. A. CULP, G. W. KINDER and
F. H. MYERS**

(Manuscript received July 20, 1970)

The single memory system provided in the Stored Program Control (SPC) No. 1A is built around the piggyback twistor (PBT). Each SPC 1A store provides 770,048 bits of random-access nondestructive readout (NDRO), electronically alterable information organized into 16,384 words. The store has a read-cycle time of 6.3 μ s.

The paper describes in detail the memory medium, memory circuits, and equipment design. A general description is given of store operations and store diagnostic and maintenance procedures.

I. INTRODUCTION

Digital storage in the Stored Program Control (SPC) 1A is provided by the piggyback twistor (PBT) memory store. The heart of this system is the PBT module which is composed of pairs of wires wrapped with two magnetic tapes, one piggyback on top of the other. The SPC 1A store has been designed around the PBT memory module. This store meets the SPC 1A requirements of high reliability and low cost with an operating speed that is adequate for the present SPC 1A applications. High reliability is provided by the nondestructive readout (NDRO) characteristic of the PBT memory. The memory element needs no regeneration after reading and is not disturbed by power interruptions. It is alterable at electronic speeds with this feature being protected from accidental use by several independent checks on the write process. Low cost has been achieved by high-volume continuous fabrication of the PBT memory element.

The SPC 1A store¹ is used for all the memory needs in the SPC 1A. This is in contrast to No. 1 Electronic Switching System (No. 1 ESS) in which the memory needs were split into two groups, temporary memory for the storage of call-related data and semipermanent

memory for storing programs and translation data. The temporary memory, known as the call store² uses the ferrite sheet as the basic memory medium. The semipermanent memory is called the program store³ and uses the permanent magnet twistor (PMT).

Figure 1 shows a photograph of a SPC 1A store frame and lists some of its important characteristics. The 47-bit word consists of 40




Fig. 1—Photograph of SPC 1A Store and Table of important characteristics.

bits of information, six Hamming check bits, and one parity bit. This provides single- and double-error detection over the word and its address as well as single-error correction on the information. The capacity of 16,384 words allows SPC No. 1A storage to be provided economically for applications that require moderate amounts of program and data.

The store read-access time is under $3.25\ \mu\text{s}$ with a $6.3\ \mu\text{s}$ read-cycle time. The write time for a single store is $50.4\ \mu\text{s}$, but multiple stores in a single system permit the use of an overlap write mode in which one store can begin a read operation before another store has completed a write. With this overlap feature, a write instruction requires between $18.9\ \mu\text{s}$ and $56.7\ \mu\text{s}$.

In addition to the PBT memory modules, the store contains all the circuitry required to operate the memory. Access circuits select the proper memory word and apply word-line drive currents for reading and writing. Bit drivers provide information-dependent bit current for writing. Sense circuits amplify the memory output and make the one or zero determination. Circuits are included to provide a communication link with the processor for normal operations and maintenance.

Maintenance is an important consideration in the store design. Both software and hardware facilities have been provided for detecting and isolating store faults.

II. MEMORY

2.1 *General*

Of the few known approaches to a memory device that offered the features of both electronic writability and nondestructive readout, only an adaptation of the permanent magnet twistor memory—called the Piggyback Twistor (PBT) memory⁴—offered the economics that were consistent with large-capacity stores for electronic switching systems.

2.2 *Organization*

2.2.1 *Capacity*

The 15B PBT memory module has a capacity of 4096 words and 47 bits per word for a total of 192,512 bits. An SPC 1A store uses four of these modules to obtain a total of 16,384 words. Since the

PBT module is an integrated approach to magnetic memory design; that is, 4096 memory sites on a continuous pair of wires, its manufacturing yield is very sensitive to isolated defects. One defective memory site renders the entire 4096 sites on a bit line unusable. Manufacturing yields close to 100 percent are maintained for the PBT by assembling the module with additional capacity so that seven spare bit lines are available to replace any bit line having one or more sites that fail to pass the final electrical tests.

2.2.2 *Physical Arrangement*

Figure 2 is a photograph showing a 15B PBT memory module as a background for a few unfolded memory planes. The inserts are enlargements of sections of the planes. The entire memory consists of 64 of these memory planes. Each plane contains 64 words at 0.175" spacings and a plane-to-plane spacing of 0.20 inch making the 64×64 array of words nearly square. The bits are spaced on 0.05" centers along the word line. The active memory bit density is 570 bits per cubic inch. The overall module dimensions are 15" \times 15.25" \times 5.75" and the unit weighs about 27 pounds. A toroidal ferrite core per word is included within the memory structure to perform part of the address decoding function (insert of Fig. 2).

Each bit line consists of a pair of piggyback twistor wires which are fabricated by wrapping two thin, flat tapes of magnetic material spirally on a 40-gauge copper wire. Figure 3 shows a photograph of a short section of twistor wire taken on a scanning electron microscope at 500X magnification. The second magnetic tape, which is wrapped underneath the first tape, can just barely be seen. The double wrap suggested the name piggyback twistor.

Each wire of the bit line is continuous throughout the memory. The 4096 bits per line require only four connections, one on each end of the pair of wires. The actual length is about 76 feet. The 108 wires that make up the 54 bit lines (47 active and 7 spare) in the PBT module are encapsulated in plastic before module assembly. In this form, the bit lines can be handled easily and economically. The word lines, which are mounted in groups of 64, are assembled as single turn solenoids complete with their access and shuttle suppression cores. Sixty-four word line subassemblies or "planes" are bonded to the bit cable. Then the bit line cable is folded so that the planes lie on top of one another. The resulting stack of 64 planes is mounted in a lightweight metal frame.




Fig. 2—Photograph of a 15B Memory module in the background and an unfolded section of a module in the foreground showing the 64-word plane assembly and the continuous bit line cable. Inserts are expanded views of the twistor wire, access core, and shuttle suppression core.

2.2.3 Electrical Configuration

The PBT memory, like any other memory, requires drive currents to read and write information, and contains sense lines along which it transmits the information to the external readout circuits. Figure 4 shows a block diagram arrangement of the PBT.

The access drive currents are generated external to the memory. They supply the excitation to an access core matrix which is part of




Fig. 3—A photomicrograph of a piggyback twistor wire. The memory tape (AuFeCo) is on top wrapped around a #40 AWG copper core wire. The edge of the sense tape (4-79 Mo-Permalloy) can just barely be seen under the memory tape.

the memory and which converts the access drive into a current on a single memory word line. Between the access core matrix and the word lines is a second array of cores that compensates for the fact that the access cores are not perfect threshold devices. If the access cores were perfect, a single *X*-access and a single *Y*-access current would be transformed into a single word-line current. Actually, in addition to the desired word-line current, 126 other word lines have small currents induced in them. The shuttle suppression core array keeps these unwanted currents small enough to be tolerable.

The 4096 word lines shown in Fig. 4 intersect the 47 active bit lines. A current flowing in a single word line will drive all bit lines in parallel. The bit lines are shown to consist of pairs of twistor wires, terminated at one end by a nonreflective termination and 47 bit driver circuits; and at the other end by a diode array and 47 sense channels.

For writing information into a memory word, an *X* and *Y*-access

line is energized, thereby selecting a single word line in which useful current flows. At the same time, each bit driver generates current around the loop formed by the two wires of the bit line and the terminating diodes. The information written into each bit of the selected word is a function of the associated bit current polarity. The coincidence of the 47 bit currents and a single word line current writes the 47 bit word.

In order to read a PBT word, the desired word is again selected by a single X and Y -access current. This time, however, no bit current flows. The word current switches magnetic material on one wire of

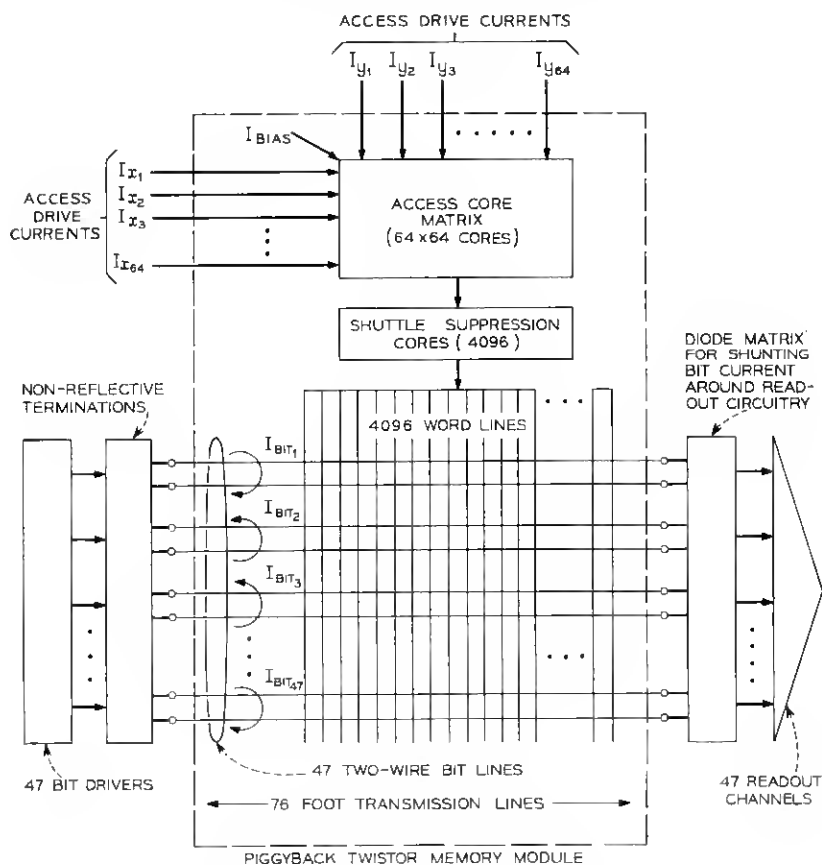


Fig. 4—Block diagram of a PBT memory showing connections to the external drive and readout circuits.

each of the 47 hit line pairs thus inducing a voltage on the bit line which now acts as a sense line. This information—either a positive or negative voltage—is transmitted down the sense/bit line to the external readout circuits.

The details of how each of the major components of the PBT contributes to the overall performance of the memory is given in the following sections.

2.3 Theory of Operation

2.3.1 Two Material System

The memory element of the PBT is made up of two short sections of twistor wire, a sample of which is shown in Fig. 3. The two tapes are mechanically wrapped on the copper wire, one to perform the information-storage function and one to sense the information. Both are made from square loop magnetic materials; that is, their characteristic hysteresis loops have vertical sides and very little slope on the top and bottom. Such loops are typical of magnetic memory materials, but the hysteresis loop of a PBT wire is not typical in two important aspects. First, it is the composite of the loops of two magnetic materials having different values of coercivity and magnetization. This produces steps in the hysteresis loop. The second difference is that there is a significant demagnetizing field caused by the fact that only a short length of wire is magnetized in operation. This demagnetizing field has the effect of shearing the otherwise vertical sides of the hysteresis loop. Figure 5 is an idealized sketch of the hysteresis loop of a PBT wire.

Of the two tapes on the PBT wires, the outside tape stores the information and the inside tape senses it. The information in the storage tape is determined by the direction of magnetization. The magnetization produces a magnetic field external to the tape that is related to the demagnetizing field. This external field can be approximated by:

$$H_{\text{ext}} = \pm \frac{2\varphi_r}{\pi l^2} : . \quad (1)$$

The sign of the external field (H_{ext}) is determined by the direction of magnetization of the remanent flux (φ_r) and the amplitude is directly proportional to the flux level and inversely proportional to the square of the length (l) of the magnetized section.




Fig. 5—The composite hysteresis loop of a PBT memory element.

2.3.2 Writing a PBT Bit

In order to write information into the PBT, a short length of twistor wire must be definable as a bit location, and the magnetization of the memory material along that short length must be driven to positive or negative saturation.

The bit location is defined by the intersection between a word line and a wire pair that makes up a bit line. The magnetization of the memory material is driven to saturation in this region by the coincidence of a word-line current and bit-line current. (Since the magnetic material is wrapped at about a 45° angle on the copper core wire as shown in Fig. 2, it is coupled by both bit- and word-line currents even though these currents are orthogonal to each other.) Figure 6 shows

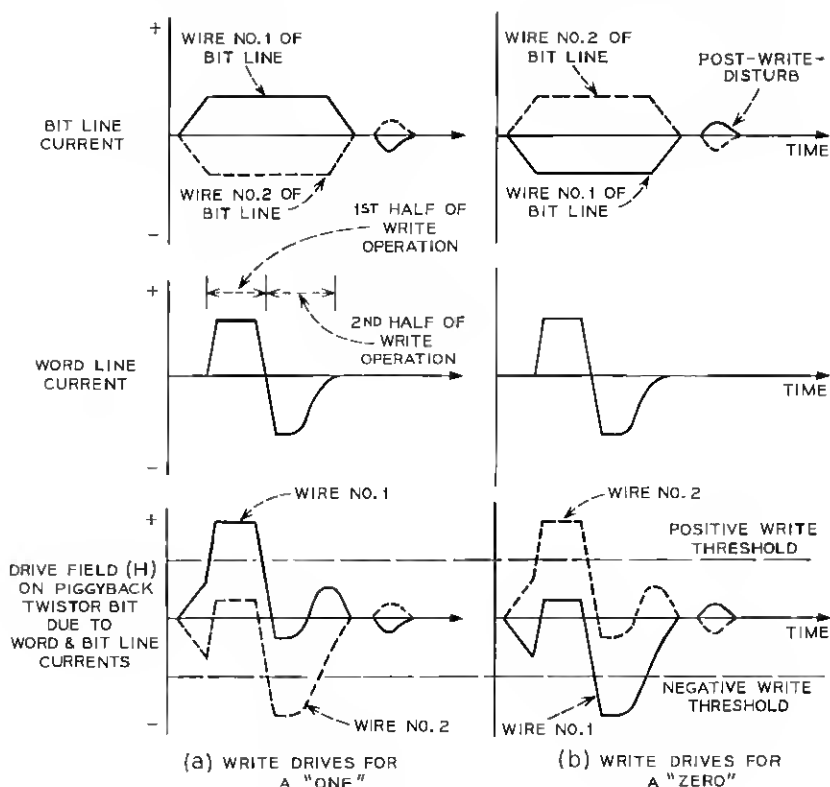


Fig. 6—Write fields acting on each wire of the bit line for writing: (a) a ONE, and (b) a ZERO.

the currents and resulting fields for writing a one or a zero. The hit current flows in opposite directions in each wire of the pair. The word-line current for writing flows in one direction for half of the write time and the other direction for the other half. For the first half of the write operation, the fields due to the hit and word current will add at one wire and subtract at the other. During the second half of the operation, the two fields will add at the other wire of the pair but with the opposite polarity than the previous summation. The result is that the magnetization of the memory material will be driven to saturation on both wires of the pair, but one will be driven positively and the other negatively. For writing information of the opposite polarity, only the bit current polarity is changed, but this

interchanges the effects on the two wires of the bit pair, reversing the magnetization in each.

For proper writing, the currents involved must be held within certain limits. The derivation of these limits will be discussed in Section 2.5. Another requirement is that the timing of the currents must be such that the bit current is maintained at its proper level during the application of the word current. When the bit current is removed, the bit line voltage is reversed momentarily resulting in a post-write-disturb (PWD) current that completes the writing operation. The purpose of this pulse is to counteract the disturbing effects of the bit current.

To insure that the bit length is the same for both polarities of magnetization (i.e., for both ones and zeros), the memory material between adjacent bits must be close to zero magnetization. If it were not, it would tend to make a bit location, of the same magnetization direction as the interbit material, appear longer than a bit location of the opposite magnetization. This would result in an asymmetrical external field [equation (1)] and an asymmetrical composite loop (Fig. 5). To avoid problems due to lack of symmetry, all bit lines are demagnetized during memory manufacture. During memory use, no drive fields are applied to the bit lines with amplitudes sufficient to disturb the zero magnetization level between bit locations. Only the individual bit locations are driven to saturation.

2.3.3 *Reading a PBT Bit*

The reading of stored information depends on the sense material magnetization state being a function of the memory material state. This is accomplished by the action of the external field [equation (1)]. The amplitude of this field is sufficient to bias the sense material beyond its switching threshold. The polarity of this field, which is determined by the magnetization direction of the memory material, either permits or prevents sense material switching as a result of a word-line current. For example, in Fig. 5, for a read pulse as shown, the sense tape would be switched if it were under the influence of a positive bias, and it would not switch under the influence of a negative bias. The read pulses must not exceed the disturb threshold or else the memory material flux will be switched, thus altering the stored information.

For the two-wire bit, information is written so that there is one and only one wire biased to switch. The information—whether a one or

a zero—is determined by which of the two wires is biased to switch and, consequently, whether a positive or negative readout voltage is induced on the bit line.

Reading, then, is accomplished by impressing a current on a word line that had been used in the past to write the desired information into the 47 associated bits. This current results in switching sense material on one of the wires of the bit pair and 47 voltages—some negative, some positive, depending on the stored information—are induced on the bit lines and transmitted to the input of the sense circuits. After the read current is removed from the word line, the external field from the memory material resets the sense material making the read operation a nondestructive one. An unlimited number of read operations are possible on any memory word without any intervening write operations.

2.3.4 *Access Core*

The access core switch is similar in operation to the access cores used in permanent magnet twistor memories (PMT). It performs part of the memory address decoding function by virtue of its threshold characteristic. Each word line in the memory is coupled to an individual access core. The cores are arranged in a 64×64 array. Each column of 64 cores has a two-turn winding threading the cores, and each row of 64 cores has another set of two-turn windings. A single-turn bias winding threads all 4096 cores in the array. The bias winding is energized with dc to bias all the cores to saturation. In order to select a word line for reading or writing, the row and column windings, which thread the core associated with the proper word line, are energized with a current pulse. The combination of these pulses is enough to overcome the bias on that particular core, switching it and inducing a current on the word line. The other cores on the selected row and column are not switched because they receive only a single pulse, which is not sufficient to overcome the core bias. This results in a linear select mode of operation for the PBT. Using a ferrite core as the access device, instead of the more commonly used semiconductor diode, decreases the number of access switches from that required for diode access, provides a current step-up to the word line by transformer action, and also provides a bipolar word-line current by the resetting action, thereby eliminating the need for a bipolar current driver.

There are two major differences in the use of the access core in the

PBT compared to its use in the PMT. The first is that the core is used to write information into the PBT. This means that significantly higher currents must be induced in the word line and, because of the two-wire bit, these currents are bipolar as shown in Fig. 6. The second difference is that the read output is sensitive to the current induced in the word line when the access core is being reset by its bias. Unless this current is limited properly, the timing of the output signal can change, the read-cycle time can be affected, and the NDRO feature can be destroyed.

In order to perform the write operation, the access core must contain sufficient magnetic flux to induce currents in the range of 3.8 amperes in the word line for a duration of about 6 microseconds. As long as the *X*- and *Y*-access pulses are of sufficient amplitude and duration, the write pulse will be bipolar—the bias current resetting the core will induce the opposite polarity word current. The duration of the reset current is determined by the amount of access core flux that had been switched by the coincidence of the *X* and *Y* pulses. If these pulses are too short, the core flux available on reset may be insufficient to write the second wire of the bit pair. For optimum writing the bipolar currents on the word line should be symmetrical.

During the read operation, the opposite problem occurs. The read is a linear select operation; i.e., all 47 bits of a word are read out in parallel by a single current on the word line. This current is induced by switching the access core in the same way that the write current was induced. In this case, however, the appropriate word line current is in the range of 1.7 amperes. Too large a current during the read operation will approximate a write-access current. This will not write information because of the absence of bit current but it will destroy stored information. This effect can be prevented on the initial read pulses by limiting the amplitude of the *X* and *Y* pulses. In order to limit the reset current, however, the amount of core flux switched must be limited. This is accomplished by a maximum limit on the duration of the *X* and *Y* read pulse.

2.3.5 *Transmission Line Effects*

Since each bit line has a capacity of 4096 bits, its physical length is relatively long—approximately 76 feet. This length and the frequency spectrum of the output signals from the bit locations make the transmission line characteristics of the bit line important memory parameters. Of particular interest are the velocity of propagation, the

attenuation, and the phase shift. The effect of these characteristics on the memory output can be seen in Fig. 7. This figure shows an oscillogram of two output-voltage wave forms from a PBT memory—one from an address adjacent to the readout terminals (near end) and the other from an address 4095 bits away from the readout terminals (far end). The oscillogram shows the far-end signal on the right to be lower in amplitude, longer in duration, and delayed when compared with the near-end signal. The variability of the output signal due to transmission line characteristics increases the difficulty of signal discrimination. However, because the transmission delay is a function of address, discrimination circuits can be designed to compensate for the variable delay of the memory output.

2.4 Design Details

2.4.1 Magnetic Materials

The choice of magnetic materials for the two twistor tapes that are wrapped on the twistor wire is based on a compromise among design parameters. The sense tape requires a relatively low switching threshold so that it can be biased easily by the memory tape and also so that it can be switched with relatively low drive currents. The total flux of the sense tape determines the amplitude and duration of the output signal and should therefore be as large as possible. However, total flux is restricted by the limitation on tape cross-sectional area necessary to




Fig. 7—Output waveforms showing the effect of bit position on the bit line. The earlier, higher amplitude signal is at the near end and the later, lower amplitude signal is at the far end with reference to the read amplifier.

minimize eddy currents, which tend to impede flux switching. In addition, the sense tape flux must be considerably less than the memory tape total flux so that a sufficient bias can be developed by the memory tape. The coercive forces of the two materials should be in the ratio of about 15 to 1 with the memory tape being higher so that it will not be disturbed by read currents. Minimum read cycle time and acceptable drive current variations also require trade-offs in material characteristics. The switching time produced by the read current is expressed by:

$$\tau_{\text{read}} = \frac{S_w}{H_{\text{read}} - (H_0 + H_{\text{ext}})}$$

where:

τ_{read} = time required to switch 90 percent of the tape flux

S_w = switching coefficient of the tape

H_{read} = read current field

H_0 = switching threshold of the tape

H_{ext} = external field from memory tape [see equation (1)].

The time required for the external field from the memory tape to reset the sense tape is:

$$\tau_{\text{reset}} = \frac{S_w}{(H_{\text{ext}} + H_{\text{read}}) - H_0} \quad (3)$$

Because H_{read} is a function of time, its polarity and amplitude in equation (3) differ from those in equation (2). Since the total memory read cycle is related to the sum of τ_{read} and τ_{reset} , there is an advantage to minimizing this sum. This requires a compromise value for H_{ext} . Equation (1) shows that the value of H_{ext} is determined by the memory tape properties. A limitation on H_{ext} that is not expressed in equation (1) comes from the fact that H_{ext} affects the memory material as a demagnetizing field; that is, it tends to reduce its magnetization to zero. In order to assure a stable state of magnetization in the memory material and thereby insure against a loss of information, the switching threshold of the memory tape (H_{0m}) must be considerably larger than H_{ext} . That is:

$$H_{0m} > H_{\text{ext}} \quad (4)$$

The value of H_{0m} cannot be increased without increasing the write current amplitudes.

The relationships that determine the field intensities required for proper writing are:

$$H_{\text{word}} + H_{\text{digit}} \geq H_{0m} + H_{\text{ext}} \quad (5)$$

$$H_{\text{word}} - H_{\text{digit}} \leq H_{0m} - H_{\text{ext}} \quad (6)$$

$$H_{\text{bit}} \leq H_{0m} - H_{\text{ext}} \quad (7)$$

where

H_{word} = field intensity of the word line write current

H_{bit} = field intensity of the bit current.

All these inequalities must be satisfied concurrently in order to have proper memory writing. In order to minimize write currents on both the word line and bit line, H_{0m} must be minimized.

A compromise among the requirements expressed in equations (1) through (7) coupled with a desire to maximize the amount of signal induced by switching the sense tape, led to the selection of material properties. A 4-79 molybdenum permalloy was selected for the sense material. It is processed to give a coercive force of 0.75 oersted and a total flux of approximately 1 mV- μ s. The memory material is a 4Au-12Fe-84Co alloy⁵ with a 13 oersted coercive force and a total flux of about 5 mV- μ s. Besides the fact that these two materials can be processed to produce compatible properties, a second outstanding characteristic of these materials is the lack of strain sensitivity. Thus, inexpensive assembly techniques can be used in memory production.

2.4.2 Ferrite Cores

Each word line in the PBT memory threads two ferrite cores—an access core and a shuttle suppression core. (See the inserts of Fig. 2.) The access core uses a square-loop material so that it acts as an address decoder. The core has low loss so that it is an efficient current transformer, and has sufficient flux to sustain the high write currents for a sufficient duration. The core that was selected was a Cadmium ferrite with a nominal cross section of 0.023 cm² and a mean diameter of 0.314 cm. This produces a loss of 0.3 ampere-turn and a total flux greater than 900 mV- μ s.

By itself, the access core is acceptable for driving the selected word address—its flux is sufficient and its loss is low enough. However, it is not an adequate decoder because it produces a shuttle current when it is one of the 63 cores on a selected row or column that receives a single pulse. Outputs induced by 126 of these currents can obscure the output from the selected core. To remedy this situation, the additional shuttle suppression core is used. This core adds a large induc-

tance in series with the word line impedance, thereby reducing the word line current induced by a half-selected core. (See Fig. 8a.) In order that the fully-selected word line current is not reduced by the inductance of the shuttle suppression core, this core is designed to saturate whenever the word line current approaches a full select current. (See Fig. 8b.) The added nonlinear inductance minimizes noise currents without seriously affecting drive currents. To do this, the core is made from a linear ferrite with a total flux of 100 mV- μ s and an ID of 0.038 inch.

2.4.3 Plastics and Adhesives

The low assembly cost of the PBT memory is due in large measure to the generous use of plastics and adhesives. The 108 twistor wires




Fig. 8—Word line current waveforms showing: (a) Current on a half-selected word line with and without a shuttle suppression core. (b) Current on a fully-selected word line with and without a shuttle suppression core.

which form the bit line cable are economically assembled into the memory by first encapsulating them in a polyester sandwich. Each side of the sandwich is made up of a 1-1/2 mil polyester film coated on one side with 2-1/2 mils of polyester adhesive. The adhesive sides of the two films are brought together with the wires between them and bonded under heat.

The materials and process for this assembly had to be designed to produce a structure that has long-term stability under operating conditions that could include short-term temperatures as high as 180°F within the memory. The use of preshrunk film and high temperature adhesive, rigid control of encapsulating temperatures, and film tensions offer the solution. The success of this assembly technique is greatly enhanced by the use of strain insensitive magnetic tapes on the twistor wire.

A similar encapsulating process is used to assemble word lines in groups of 32. A single film is used in this case. One side of the flat copper word lines is embedded in the adhesive. Long lengths of word line cable are produced and cut into short lengths to make the word line loop. Two of these 32 word line subassemblies are required per memory plane.

A specially developed pressure sensitive adhesive is used to bond the plastic encapsulated bit line cable to the word line subassemblies. Experience showed that using only a pressure-sensitive adhesive for this bond would not result in a uniform and stable spacing between the word line and bit line and this dimension critically affects the transmission line properties of the bit line. In order to stabilize this dimension, a 1/2 mil plastic film coated on both sides with pressure-sensitive adhesive is used. The film adequately stabilizes the spacing between the word lines and bit lines. A cross section of the word line and bit line cable assembly is shown in Fig. 9.




Fig. 9—Cross section of a bit line and word line assembly.

2.5 Operating Characteristics

2.5.1 Drive Variations

Since selection of the magnetic material properties determines the operating currents, any practical design must assume operation over some variation in drive currents. The PBT is designed for access drive current variations (X -axis, Y -axis and bias) of ± 2 percent. These currents determine word line current according to the equation:

$$I_{\text{word}} = 2(I_x + I_y) - I_{\text{bias}} - I_{\text{loss}} \quad (8)$$

where I_{loss} is a loss term due to the energy required to switch the access core. A study of equation (8) will show that the 2 percent variation in access currents amounts to a considerably larger variation in word line current. Table I summarizes the drive variations throughout which the PBT memory cell must operate.

2.5.2 Output Signal

The output signal from the PBT is bipolar. The polarity is related to the polarity of bit current that was used to write the information. The output amplitude ranges from a minimum peak of 3.5 mV to a maximum of over 10 mV depending on the drive levels, information stored, operating history, etc. The width of the output signal at the 0.5 mV level can have a duration of about 900 ns but this duration is decreased and its timing is varied by the operating conditions. Figure 10 shows the effects of some of the operating variables on the output signals. Figure 10a is a photograph of 64 ONE outputs from a single

TABLE I—DRIVE VARIATIONS

	Maximum (amps)	Minimum (amps)	Variation (%)
Bias I_{bias}	3.98	3.82	2.0
Read Access			
I_x	1.48	1.42	2.0
I_y	1.48	1.42	2.0
I_{word}	1.8	1.4	12.5
Write Access			
I_x	2.04	1.96	2.0
I_y	2.04	1.96	2.0
I_{word}	4.04	3.56	7.7
I_{digit}	0.15	0.21	16.6




Fig. 10—ONE outputs from a PBT: (a) 64 outputs from a middle plane with noise canceling pattern and no read pulse amplitude variation. (b) 128 outputs from near end and far end planes with maximum noise pattern and extreme read pulse amplitude variations.

memory plane. Since the outputs are from the same plane, there is very little variation due to transmission line effects. All the bits are read with a constant amplitude read pulse and a minimum noise information pattern has been stored in the rest of the bit line. By contrast, Figure 10b shows 128 ONE outputs—64 from a near end plane and 64 from a far end plane. The bits are read with both

extremes of read pulse amplitude and the information pattern ends to maximize the negatively going noise output. A comparison of the two photographs shows how the area between the bundle of output traces and the zero line decreases as the number of operating variables is increased.

2.5.3 *Transmission Delay*

The average transmission delay of an output signal at room temperature is about 10.6 ns per memory plane. This varies about ± 6 percent with information content. A bit line having all ONES or all ZEROS stored will have a total delay of about 640 ns. That same bit line with alternating ONES and ZEROS on adjacent bits will have a total delay of about 720 ns. This change in delay is due to the fact that bit line inductance is related to the information content of the bit line.

2.5.4 *Temperature Effects*

Magnetic metals with Curie temperatures of 700°C to 800°C, well in excess of any practical operating temperature, from the basic memory element in the PBT memory and, therefore, the characteristics of the memory cell show no temperature effects. However, both access and shuttle suppression cores do show temperature effects. The transmission line characteristics of the bit lines change slightly due to temperature because of the conductivity changes in the copper core wire and the dimensional variations in the plastics. All these changes are small. The only temperature effect that is important enough to require compensation is the output signal delay. As the temperature increases, the output is generated earlier in time because of the changing characteristics of the ferrite access core. This change is nominally 1.0 ns per °C.

2.6 *Worst-Case Testing*

In a memory that has as many operating variables as the PBT, an extensive test sequence is needed to guarantee all the requirements implied by random access, electronic writability and nondestructive readout performance. This is accomplished by a sophisticated test machine that uses a small digital computer for address and function control, and pulse measuring and recording equipment for data display. The test sequence is a combination of worst-case information patterns, including history effects, and extreme values of drive cur-

rents. Testing at temperature extremes is performed on a sample basis. Outputs under all these variations are tested for proper polarity, amplitude, and duration. The output noise level prior to the signal interval is also tested as an additional check on memory quality. Figure 11 shows a block diagram of the PBT test facility.

III. ACCESS

3.1 General

The access system must provide the regulated read or write current drive pulses to the PBT modules. It also provides the regulated DC bias for the biased core access matrix. The access method and many of the circuit designs are similar to those used in the permanent magnet twistor³ program store of ESS No. 1. Circuits were modified wherever necessary to accommodate the higher current levels and the different timing of the PBT module.

3.2 Operation

Each PBT module has 64 *X*- and 64 *Y*-access lines. Four PBT modules are arranged in a two-by-two matrix with *X* or *Y* lines from adjacent modules connected in series as shown in Fig. 12. The result is a 128 by 128 array of access wires. To select a desired word, a current is applied to the proper *X* and *Y* wire.

Figure 13 shows the general arrangement of the access system for

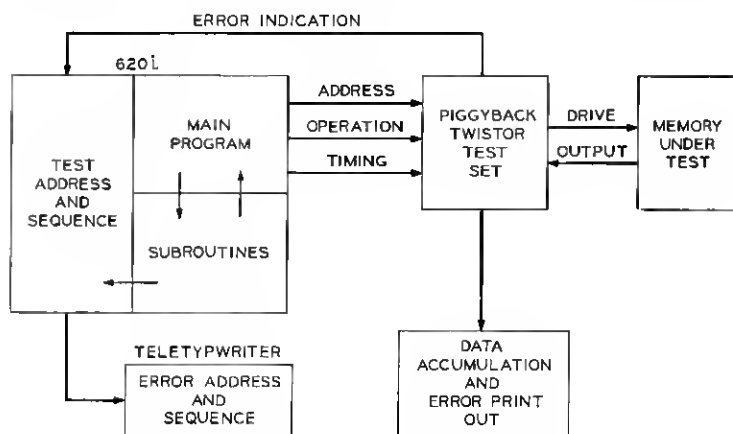


Fig. 11—PBT test facility block diagram.




Fig. 12—Module access arrangement.

one axis (X or Y). This arrangement is duplicated for the other axis. One bias current regulator, which is not shown, is provided for the whole store. The 128 wires are arranged in an 8 by 16 array. The particular access wire is then chosen by selecting one out of eight lower access switches and one out of sixteen upper access switches.

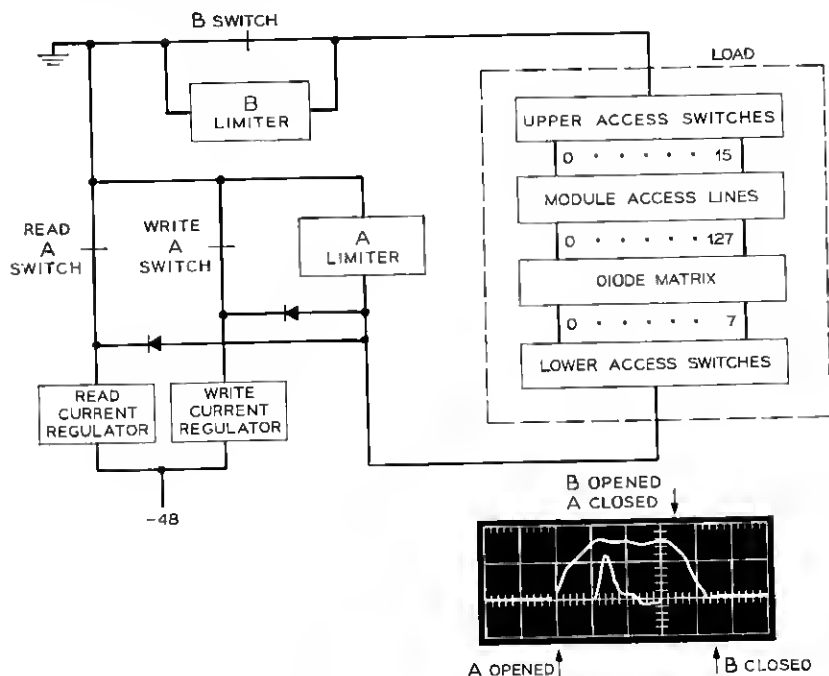


Fig. 13—Access system for one axis.

Three bits of the address are used on each axis for lower switch selection and four bits for upper switch selection.

The regulated read and write currents normally flow to ground through their associated *A* switch. When the module is to be accessed, the address information is used to first turn on the access matrix switches associated with the selected address. After allowing sufficient time for the matrix switches to operate, the desired (read or write) *A* switch is pulsed open, forcing the current to flow through the associated steering diode, the access matrix switches, and the *B* switch to ground.

The *A* limiter clamps a constant voltage across the access system during access rise time. This forces a current ramp through the primarily inductive access matrix. The access current waveform is shown in Fig. 13. When the current reaches the flat top level, the *A* limiter ceases conduction.

One microsecond later, the *A* switch is closed and the *B* switch is opened. This initiates the fall time interval. The *B* limiter clamps the matrix voltage to set the fall time. After the matrix current reaches zero, the *B* switch is closed and the access matrix switches are opened.

3.3 Current Regulators

The active series current regulators represent a major departure from the switching mode regulators used for the PMT. Figure 14 shows a partial schematic of the regulator.

The circuit functions by comparing the drop across *R*₁ to a regu-




Fig. 14—Access current regulator.

lated voltage and adjusting the drive to the Darlington configuration of transistors Q_1 , Q_2 , and Q_3 to maintain the voltage equal to the reference. The reference is an aged zener diode (CR1) driven by a constant current source. The "output adjust" potentiometer provides some adjustment range to compensate for initial tolerance on CR1 and R1, and offset in the differential amplifier. The reference resistor (R1) is a precision power resistor chosen for good initial tolerance and tight load-life tolerance.

The differential amplifier is a matched transistor pair, with a constant current load. The series regulator uses a triple Darlington circuit to minimize base drive effect on output current. Both Q_2 and Q_3 are mounted on heat sinks to accommodate the high power dissipation. Resistor R2 shunting the series transistor provides most of the output current, thus lowering dissipation in the series transistor.

The regulator loop is stabilized by a single capacitor across the differential amplifier. Test points TP1 and TP2 are provided to allow a field check of the regulating voltage (nominally 6.00 volts). The inductor (L) absorbs the module back voltage during access current pulsing, or, on the bias regulator, provides the high bias source impedance needed by the access matrix.

This same design is used for the read, write, and bias regulators, with only the values of R1 and R2 changed to provide the currents shown in Table I.

The flat top regulation is determined primarily by the regulation of X , Y , and bias regulators. End-of-life regulation is ± 2 percent in each case. Regulation is primarily controlled by the current sensing resistor in the regulator (R1 in Fig. 14). Measured short term regulation with power supply (42.7 to 52.5 volts) and ambient temperature (2 to 50°C) variation is less than ± 0.35 percent. The ± 2 percent tolerance of the individual regulators is magnified by the access matrix, resulting in a ± 12.5 percent solenoid drive variation which meets the PBT memory module requirement.

IV. READ

4.1 General Organization

In the module arrangement shown in Fig. 12, whenever a given module is accessed, the half-access current of either X or Y access alone flows through two other modules. Although shuttle suppression cores are used, the shuttle noise during readout is still significant.

The module which is diagonal to the selected module however, receives no half-access currents, and so has a quiet bit line. Therefore, the bit lines of modules 0 and 3 are connected in parallel, and those of modules 1 and 2 are connected in parallel. This always places the outputs of a quiet module and a fully-accessed module in parallel. The paralleling has the undesired effect of reducing the readout signal amplitude by a factor of two, but it does require only one-half as many readout preamplifiers, as well as simplifying preamplifier selection.

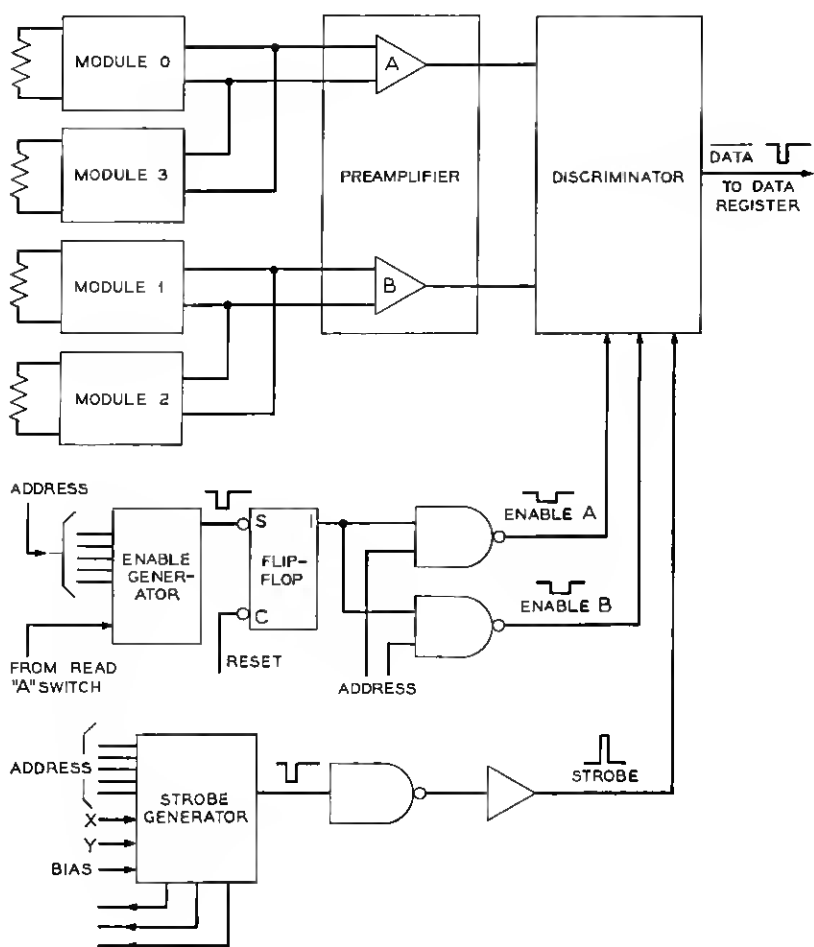


Fig. 15—Readout system.



Fig. 16—Relative readout system timing.

Figure 15 shows the block diagram of the readout system. The outputs of the two preamplifiers, *A* and *B*, are fed to the zero-crossing discriminator, where selection is made by using two separate enable signals, controlled by *X* and *Y* address bits. The enable signal is started when current through the access *A* switch is interrupted (at about 10 percent of access drive). The enable generator then introduces both a fixed and an address-dependent delay before generating its enable pulse to set the read enable flip-flop. This flip-flop is reset after strobing of the information into the data register.

The signal from the zero-crossing detector in the discriminator is gated to the data register by the strobe pulse. Since the strobe timing is more critical to proper operation of the readout system, it is generated with greater relative timing accuracy. This is accomplished by starting the strobe reference timing from a fixed point on the read access overdrive. The proper timing is derived by using *X*, *Y*, and bias currents from the access system and triggering the delay start with the optimum 1.0 ampere-turn overdrive point. The strobe generator then produces its fixed and address-dependent delays and generates the 150 ns strobe pulse. This is then amplified and fed to the discriminators. The relative timing of these signals is shown in Fig. 16.

4.2 Preamplifier

The readout amplifier raises the output of the PBT modules to a level of 500 mv sufficient for gating and polarity detection in the discriminator. It is a two-stage direct coupled amplifier with a minimum gain of 2000 and 6 dB response points at approximately 150 KHz and 1.5 MHz. The low frequency cut off is necessary to insure fast amplifier recovery after a write current pulse. In addition, a resistor diode network is used at the input to limit the overload the amplifier sees during the write process.

4.3 Discriminator

The discriminator circuit provides selection of either the *A* or the *B* preamplifier and polarity detection of the selected signal. The output for a ONE signal is a low-going logic pulse used to set the associated data register flip-flop.

Figure 17 shows a partial schematic of the discriminator. The inputs from both preamplifiers are AC coupled to diode bridge clamp circuits. The clamp circuits provide a low impedance ground for the preamplifier signals until one side is unclamped by the enable pulse, which back biases the clamping diodes in that particular bridge. Resistors R_1 , 2, 3, and 4 form a resistive adder to couple the enabled signal into the differential detector Q_1 and Q_2 . The output of Q_2 is coupled through zener diode CR_1 into the output gate, and ANDED with the strobe to produce the data output.

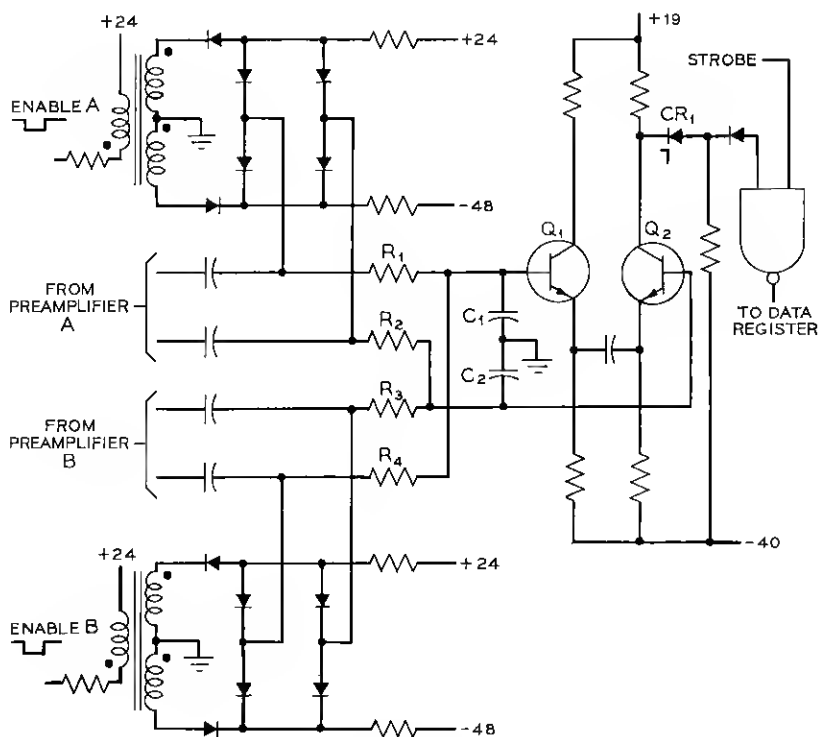


Fig. 17—Readout discriminator.

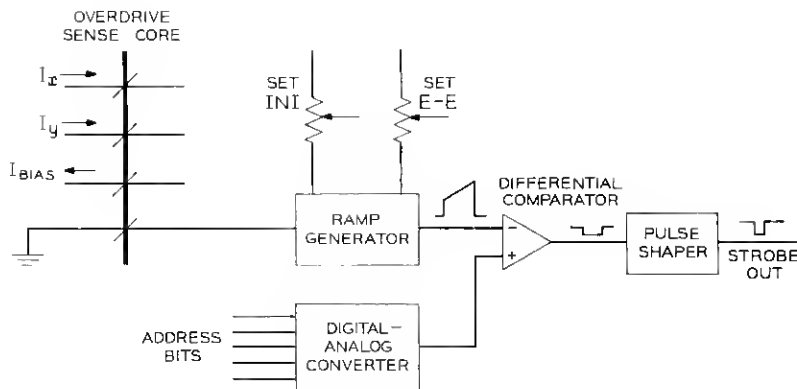


Fig. 18—Strobe generator.

The differential amplifier is greatly overdriven by the amplified output to provide zero-crossing detection. Capacitors C1 and C2 provide low-pass filter action before the differential detector. This effectively lengthens the readout signal, giving a larger operating window. Two discriminators are provided on a single circuit board, with the enable bridges being shared by both circuits.

4.4 Strobe Generation

The strobe generator and associated fanout amplifier chain must generate the address and temperature dependent strobe pulse and deliver this pulse to the 47 readout discriminators. Figure 18 shows a block diagram of the strobe generator circuit pack.

The overdrive sensing core is selected from PBT access matrix cores and threaded with the X , Y , and bias current leads. The turns are arranged so that it produces an output pulse at the 1.0 ampere-turn overdrive point. Using an access matrix core provides first-order temperature compensation for the output signal delay temperature effect described in Section 2.5.4.

The five address bits, which partially determine the plane to be accessed, are used as inputs to the digital-analog converter. The analog output then is proportional to the distance along the twistor tape from the readout amplifier, and thus proportional to the expected time the output will occur.

The pulse from the overdrive sensing core is used to start a ramp generator. The ramp generator output is then compared to the output

of the digital-analog converter, and the comparator output used to initiate the strobe pulse.

Adjustments are provided on the strobe generator board to set both the initial (address independent) and the end-to-end (address dependent) values of the delay. These are factory adjusted to compensate for component variations on the strobe board itself.

4.5 *Enable Generation*

The enable pulses for the discriminators are generated in much the same way as the strobe pulses. The accuracy of positioning is less critical, however, so the parameters are not as tightly controlled. The address tracking is accomplished in the same manner as for the strobe. The pulse must be considerably earlier than the strobe, however, so the start of the ramp is initiated by a pulse from the "A" access switch when it turns off. Starting it from overdrive would not allow it to be early enough for optimum positioning. The enable overlaps the strobe, so the pulse from the enable generator is used to set a flip-flop which is reset after the strobe interval. The flip-flop output is then used to drive high power logic gates for fanout to the discriminators.

4.6 *Performance*

Readout system performance has been excellent. The store operating margins are sufficiently large that once the strobe is adjusted the error rate is zero over all operating conditions. No failures have been encountered due to drift.

V. WRITE

5.1 *General Organization*

Figure 19 shows the circuits for writing one bit line in each module. Forty-seven of these circuits are provided in each store. The four modules are effectively connected in parallel during a write current pulse, and are isolated by the reverse biased diodes at other times. The writing of four modules in parallel is dictated by economics, and is the limiting factor on bit current regulation. The readout preamplifiers are connected across the near ends of two modules in parallel, and the far ends of all modules have resistive terminators (RT). The write current driver provides the bidirectional current switching and the current regulation.

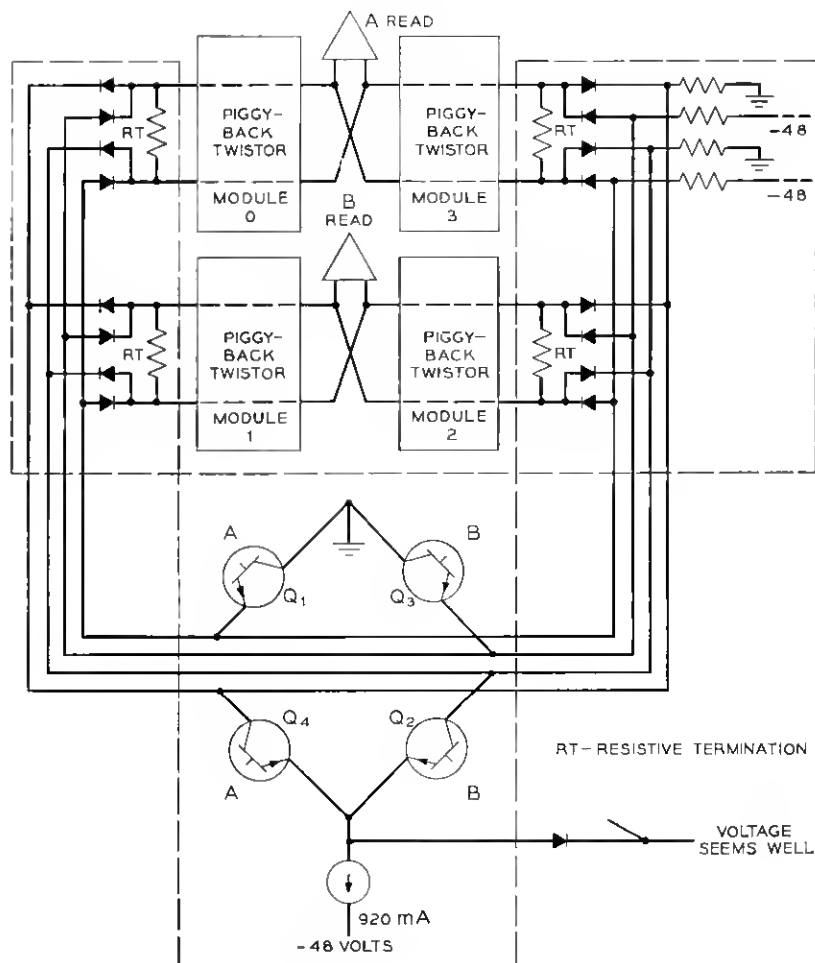


Fig. 19—Module bit interconnections.

5.2 Bit Current Driver

Figure 20 shows a simplified schematic of the bit current driver. To write a ONE, a positive write current must be followed by a negative post-write-disturb. For a ZERO, both currents must be reversed. The logic to accomplish this is done by the AND and NOR gates shown. The WP (write positive) and WN (write negative) pulses correspond to the "write" and "post-write-disturb" pulses for modules

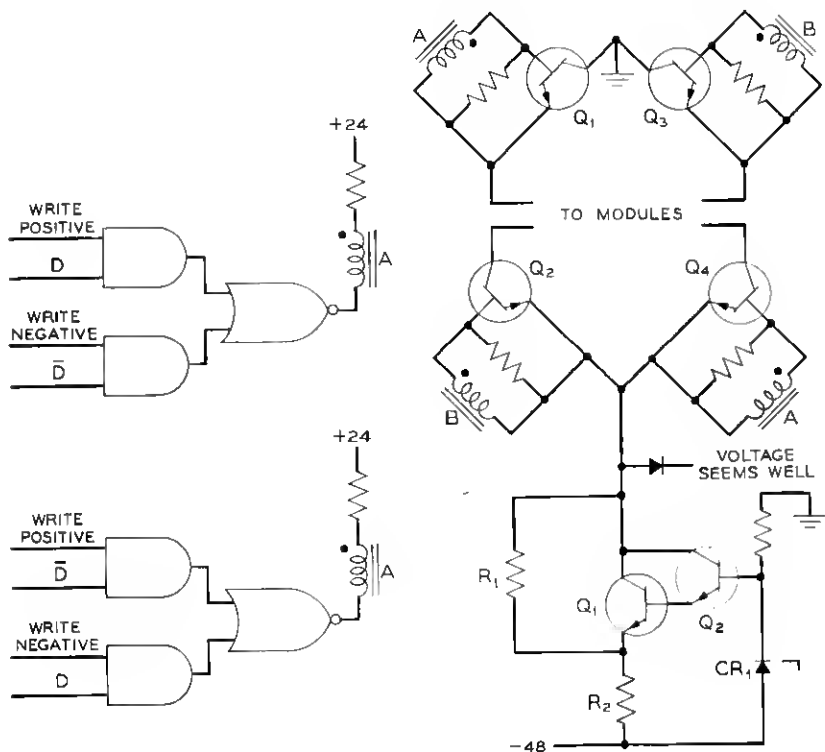


Fig. 20—Bit current driver.

0 and 1. Since modules 2 and 3 are connected in reverse to the readout amplifiers, the WP and WN functions are interchanged by logic preceding the bit current drivers during writing in these two modules.

The NOR gate outputs drive either the A or B transformers to generate drive for the proper two bridge transistors.

The current regulator is composed of the Darlington pair of Q1 and Q2, CR1, R1, and R2. The dissipation in Q1 is kept low by bypassing most of the current around Q1 by R1. One bit current driver is provided on a single circuit pack.

The voltage seems well (VSW) test point is monitored during the bit current pulse to check for shorted bridge transistors. If a bridge transistor is shorted, the VSW point will rise almost to ground, causing an ASW failure to be sent to the processor.

During the post-write-disturb interval, the current limiter is satu-

rated due to the high load inductance. At this time modules are driven by the -48 volt supply. To maintain the necessary volt-microsecond control, the length of the post-write-disturb is varied according to the potential of the -48 volt supply. The pulse generator circuit is designed to keep the volt microsecond area constant with supply voltage changes from -42 to -53 volts.

VI. READ-WRITE TESTS

The object of the read-write tests is to establish all the worst case operating conditions in the frame and then evaluate the read-write margins.

The variables that must be controlled to establish worst case read-write operating conditions are those which affect the module operation and are listed in Section II of this article. To carry out worst case testing required a facility utilizing a computer-controlled test set.⁶ This facility is shown in Fig. 21. The computer uses a complicated algorithm to establish the testing sequence which is then applied to the store. The variables under computer control are listed in Table II. Since this facility is very flexible many testing sequences have been checked that were suspected of being worst case thus evolving a very comprehensive test.

To measure the read-write margins, the readout strobe timing is varied while store outputs are checked against expected outputs. There are two timing adjustments on the strobe board; initial (INI) and end-to-end (E-E) delay. The range of these two adjustments for correct outputs is a measure of the store frame operating margins.

For read-write tests, both strobe adjustments are put under computer control and each can be set to one of 32 possible values. The CCSTEVA (computer-controlled store evaluation) program determines the strobe settings at which the store will pass all tests at all addresses. It then produces a plot of failing and passing settings with the E-E

TABLE II—COMPUTER CONTROLLED STORE VARIABLES

	Variable
	Read-Write Program
	Strobe Timing
	Enable Timing
	Access Drives
	Bit Drives
	Frame Supply Voltage
	Post-Write-Disturb



Fig. 21—Computer-controlled test facility.

FILE OUT 0

FILE 0 DESCR= 01 PROG=05 03/14/69 PROT=0
 INH= 00 RUN= 1 MOD= ALL Q.C.=0541

CCSTEVA




Fig. 22—Strobe window.

values as the abscissa and the INI values as the ordinate. Such plots are known as "strobe windows" and one is shown in Fig. 22. It was taken on a store at room temperature with all drive currents held at nominal. The read-write program was worst case and took about 120 minutes to run. Also shown is the production test requirement.

The computer-controlled test set has been used extensively in the development of the PBT store. It is also being used for production testing. Field experience has shown that a store whose strobe board is adjusted optimally with respect to the strobe window will operate without errors as long as there are no circuit faults.

TABLE III—PBT STORE COMMUNICATION LINKS

Unit	Type Signal	Wire Pairs	Description
Input From Processor	Unipolar 0.5 μ s Pulses	71 Data 5 SYNC 1 Write Go	Duplicated address bus from the processor. Each bus is shared by half of the stores in the system.
Output to Processor	Unipolar 0.5 μ s Pulses	47 Data 2 ASW 1 SYNC	Duplicated answer bus to the processors. Each bus is shared by half of the stores in the system.
Input From Central Pulse Distributor	Bipolar 0.5 μ s Pulses	2	Dedicated leads to control the trouble and protected area flip-flops.
Input From Master Control Center	DC	2	Dedicated leads to control the trouble flip-flops.
Output to Master Control Center	DC	2	Dedicated leads to indicate the state of the trouble flip-flop and the power relay.
Input From Signal Distributor	DC	8	Dedicated leads to control the several scan relays and the out-of-service lamp.
Output to Scanner	DC	56	48 scan leads shared by all stores and 8 dedicated scan leads.

VII. COMMUNICATIONS

The major communication links between the processor and the PBT store are the address and answer buses. The signals on these buses are 0.5 μ s pulses for binary ONES and no pulses for ZEROS. The address bus carries the address, control, and data information from the processor to the stores and the answer bus carries data back to the processor. The bus systems and stores are completely duplicated with each store having access to only one bus system but with each processor having access to both bus systems. All store inputs and outputs are listed in Table III.

VIII. STORE OPERATIONS

Store operations are divided into normal and maintenance. There are three normal operations, normal read, normal write, and inactive. When the store is faulty or suspected of being faulty, various main-

tenance operations are possible. Three of these, the bus-register test (BRT), control read, and control write operations are powerful tools that use the full address and answer bus facilities of the store. If necessary, all bus communications with a given store can be stopped by setting a flip-flop in that store. This can be used to silence a store which is signaling on the answer bus when it should be quiet. When this flip-flop, which is called TBL (trouble) is set, DC measurements can be made on the store by the scanner. The TBL state can be modified by setting the PORT (partial override of trouble) flip-flop to allow the store to receive address bus inputs and execute instructions but not respond. This state is used in the store update operations after a store has been out of service.

Figures 23, 24, and 25 illustrate the major circuits that are used for store operations. On all bus operations, the 24 bits specifying the store name, operation, and address are gated to the proper registers in all stores by the SYNC pulses. If the transmitted name matches the name that is wired in the decoder, the store goes active. If not, the store executes an inactive operation.

During each bus-controlled operation the ASW (all-seems-well) circuit gathers data on the operation of critical store circuits. If they all give indication of functioning properly, it then sends an ASW signal to the processor. If the signal is not returned, the processor retries the command. If it fails again, the processor forces a transfer to the maintenance programs.

8.1 *Normal Operations*

Figure 23 shows the major information flow for the normal read, normal write, and inactive operations.

8.1.1 *Normal Read*

On a normal read, the addressed location is interrogated and the resulting 47 bits are passed to the data register and then to the processor. The data is gated out sometime between 3.1 microseconds and 3.8 microseconds after receipt of the read command from the processor.

8.1.2 *Normal Write*

The normal write operation consists of a store read (called preread) and store write. The preread operation is a normal read with the exception that the name, operation, and address registers are not reset at the end of the cycle. This serves two functions: Often, only part of

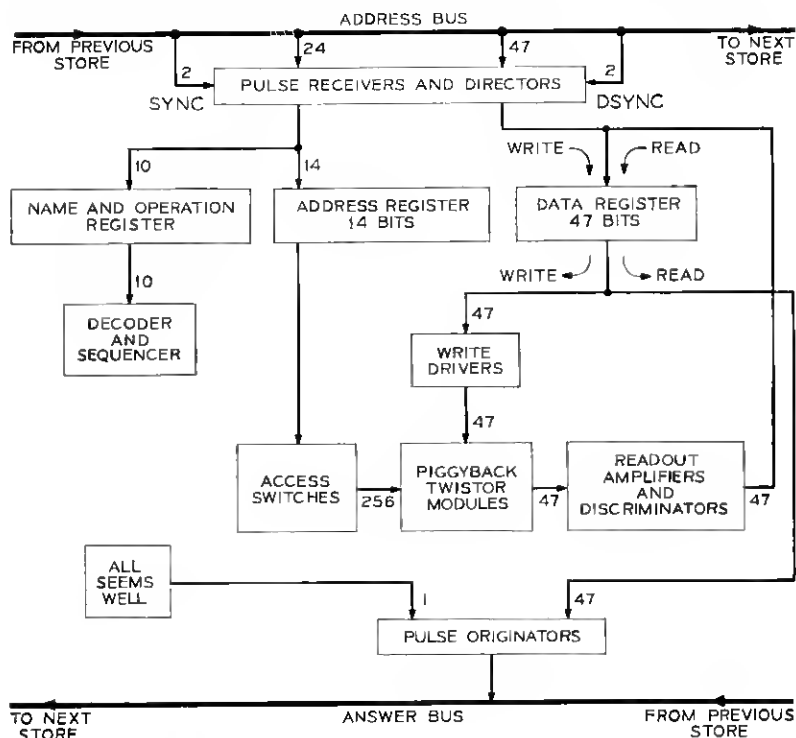


Fig. 23—Information flow for normal operations.

a word is to be modified and the preread is necessary to obtain the other bits that will not be changed since all bits must be written simultaneously.

The other preread function is that it provides a fail safe protection of the processor to store address path. The processor makes a Hamming check of the data and address of the preread cycle. If this check is passed, then the proper address has been locked into the store. Now the processor sends the 47 bits of data for the write operation which are received by the store and are held in the data register. The store writes this data at the location just read. Thus a check has been made on the whole address path from the processor to the store assuring that the correct location is written. In addition, a "write go" pulse from the processor is required during the write cycle by the store. This pulse initiates write access. The absence of this pulse causes the

cycle to abort and fail to send an ASW signal. The store is ready for another cycle 50.4 μ s after the start of the preread cycle.

The PBT store has a protected area feature that prevents programs from writing into certain preselected areas. The protected area is specified by permanent wiring in the decoder. There are seventeen possible protection options; no protection, $\frac{1}{16}$ protected, $\frac{1}{8}$ protected, etc., to all protected. All areas protected are continuous and start at address 2. Addresses 0 and 1 are always unprotected for maintenance purposes. If a write is attempted into protected area, it is recognized by the decoder on the preread operation and the store fails to go active. Since an inactive store does not respond, the processor detects an ASW failure and calls maintenance programs.

All program and translation data are located in protected areas of the store so that they will not be destroyed by programs that attempt to write in the improper area. Since maintenance programs are called by these attempts, many program bugs can be uncovered by analyzing attempts to write in protected area.

8.1.3 *Inactive*

All stores receive and register the full information for each operation. Only the store that recognizes its name will go active. The other stores perform a simple sequence. The most important part of this sequence is resetting all registers and verifying the reset with the all-seems-well circuits. Thus, if during a period of inactivity, a transient sets a register in a store, the register will be reset on the next operation in any store.

As was noted above, a write cycle takes 50.4 μ s. Much of this time is spent applying bit current and then waiting for the read circuits to recover. The processor has completed its write cycle after two cycles (12.6 μ s) and is ready to read again. Therefore, all inactive cycles are 6.3 μ s long and if the next operation after a write is in a store that is different from that just written, it will commence two cycles after the write began. This write overlap feature allows a write instruction to consume between three and nine system cycles (one cycle is added to read the write command).

8.2 *Maintenance Operations*

Store maintenance operations are used both for store fault recognition programs and store diagnostics. The store fault recognition programs, using these operations, run a very brief set of tests on a suspect

store and determine whether the store should be diagnosed or remain in service. The store diagnostics, using these operations, run a lengthy sequence of tests which will resolve a majority of the faults to a small number of circuit packs.

The communication bus is used to execute the bus-register test (BRT), control read and write operations, and to set up the trouble mode. These are all $6.3 \mu\text{s}$ operations. The signal distributor is used to set up the scan mode. Scanning proceeds at relay speeds.

8.2.1 BRT (Bus-Register Test)

The BRT operation is used for verifying the proper operation of registers, sequencers, and bus communications. The information flow for this operation is shown in Fig. 24. In this operation the six-bit name and the 14-bit address are formed into a 20-bit data word, double-gated into the data register and outpulsed to the processor. The Hamming bits cannot be checked by this method, but if this test passes, the diagnostic program can proceed. The Hamming bits will be checked in a later test.

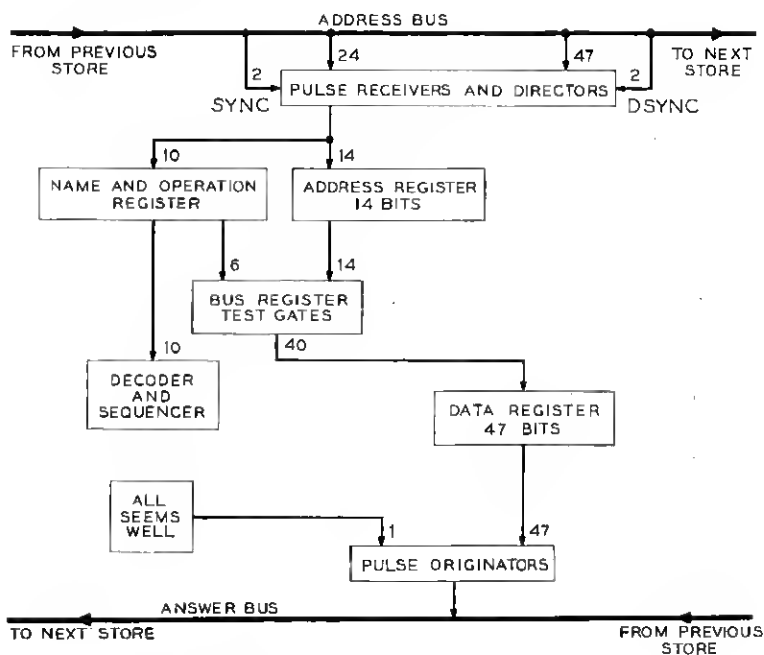


Fig. 24—Information flow for BRT operation.

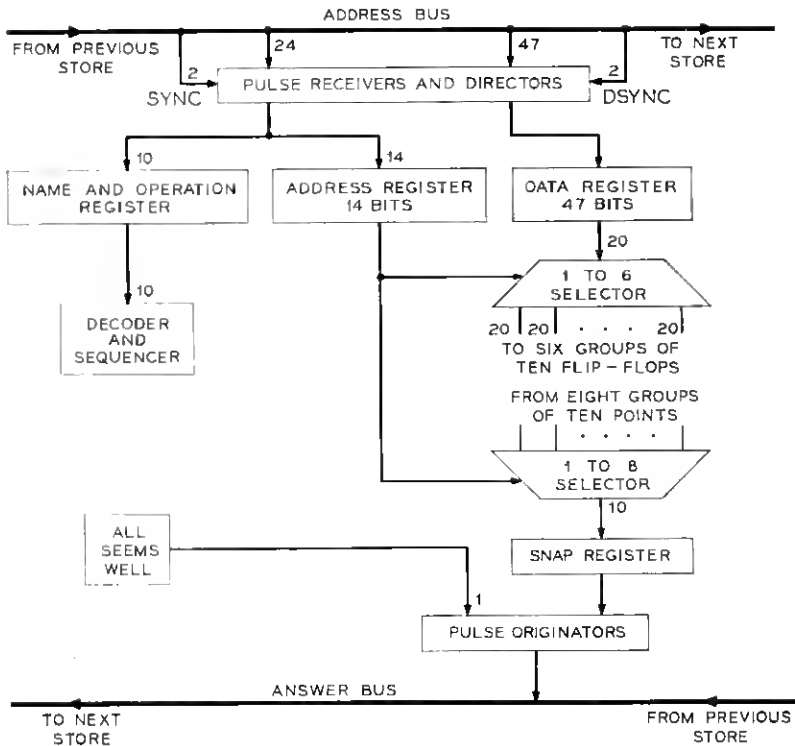


Fig. 25—Information flow for control operations.

8.2.2 Control Read

The control read operation is used by the maintenance programs to determine the state of a group of store circuits. Figure 25 shows the information paths for this operation. At one of three times in the cycle the states of a group of ten flip-flops are gated into the SNAP register. This register holds the data until it is time to send it on to the processor. Any one of eight groups of ten flip-flops may be specified. These 80 flip-flops comprise most of the control and maintenance flip-flops in the store.

8.2.3 Control Write

The control write operation is used to set the state of important store circuits. Figure 25 shows the information paths for this operation. One of six groups of twenty flip-flop gates can be pulsed in a

control write operation. After these points have been written, their states are gated into the snap register and can be read on a subsequent control read to verify the control write operation.

8.2.4 *Delayed Snap*

The delayed snap operation makes it possible to get a control reading on a normal operation by using both a control write and a control read. This is important for maintenance since many store functions are inhibited during maintenance operations.

To make a delayed snap requires a control write of the snap flip-flops to establish on which of the next seven operations the control reading is to be taken. On the specified operation, at the specified point in the cycle, the states of the selected group of ten flip-flops is gated into the SNAP register. A subsequent control read brings this information to the processor.

8.2.5 *Scan Mode*

The scan operation provides DC information about various circuits in the store. In this mode relays are operated to connect a 1.3 K Ω ferrod scan point in the scanner into the desired store circuit. The scanner is interrogated by the processor under program control and returns a ZERO if more than 3.9 mA flows and a ONE if less than 1.8 mA flows. This mode is used to test the output of the voltage regulators, the fanout diode in the timing circuits, the states of many flip-flops, and the state of the store power relays.

IX. MAINTENANCE

Every store fault must be repaired as rapidly as possible. Figure 26 shows the process used to repair stores. The next sections describe the three techniques used to identify the fault.

9.1 *Diagnostic Program*

All attempts to repair stores begin by running the diagnostic program which usually generates a trouble number. The diagnostic program usually is called automatically because of either massive operation failures, the failure of a scheduled exercise, or a fuse alarm.

Occasionally a fault will occur that causes a very low number of store operation failures (typically less than 30 in any 30 second interval). If these persist, the maintenance man will begin his diagnosis by calling the diagnostic program.

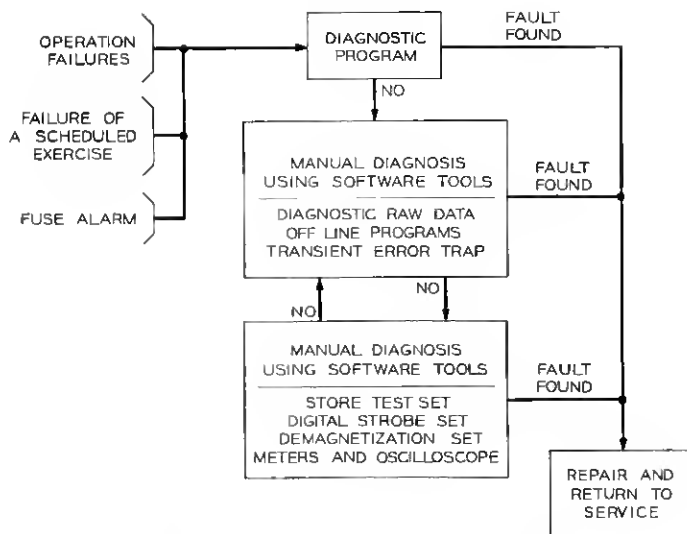


Fig. 26—PBT store maintenance process.

If the diagnostic program produces a trouble number, the maintenance man looks up the trouble number in the trouble locating manual (TLM). He usually finds the number listed and along with it is a list of circuit packs whose failure result in that number. He replaces these packs, one at a time, until the store operates properly.

Studies have shown that about two-thirds of the store faults will result in a trouble number that can be found in the TLM. Studies of the TLM show that 65 percent of these matches will identify three or less circuit packs. However, 10 percent identify more than nine circuit packs and may require manual methods to resolve the fault.

The remaining third will result in no trouble number or a random trouble number due to marginal fault conditions. In these situations manual diagnosis must be used. To date the manual diagnostic tools have proved satisfactory. As a matter of fact, several very subtle faults have been found readily and rapidly using these tools.

9.2 Manual Maintenance Using Software Tools

If a trouble number is not listed in the TLM, the maintenance man can request the diagnostic program to print out all the raw data on which the trouble number is based. This data shows exactly which

store tests failed. By analyzing the raw data he can determine areas of the store which should be checked for failure.

It may be necessary for him to cause the store to sequence to isolate the problem. To this end a set of "off-line" programs are provided that will allow him to sequence the store repetitively through as many as twenty operations. He can then use an oscilloscope to dynamically view the suspected store areas.

If the fault is a transient operational failure and the store passes the diagnostic program, it is necessary to recreate the conditions that exist when this failure occurs. To do this, programs and hardware are provided that will enable the maintenance man to pinpoint what program steps were being executed at the time of failure. He can then use the off-line program to set up this condition and to trace the fault.

9.3 *Manual Maintenance Using Hardware Tools*

When software approaches have failed to locate a store fault and the problem appears to be in the read-write circuits, it is necessary to resort to manual testing. Several pieces of equipment have been provided in each TSPS office for manual maintenance: the SPC store test set, the digital strobe set and the PBT demagnetizer set. When this equipment is used, the store is taken completely off-line and is manually disabled from communicating with the rest of the system.

The SPC store test set connects to the store via a special test connector. It gains direct access to the sequencer, decoder, read, write, access, and ASW circuits. It has facilities for sequencing several severe but not worst case read-write programs through every address in the store. On every read operation it checks the data read out and will indicate errors. Most marginal read-write or ASW faults are easily located with this test set.

The digital strobe set is used with the store test set for marginal checks on the store. By varying the strobe settings while the test set is running test patterns, a check may be made on the strobe window (see Section VI and Fig. 22). Figure 27 shows the test set and strobe set connected to a store.

If a portion or all of a bit wire is bad, there is a possibility that the memory material in the interbit regions has become magnetized due to circuit pack malfunction. To remedy this, the PBT demagnetization set is connected to the bit wire and operated. A decaying 60 Hz sine wave is applied to the bit wire which leaves it completely demagnetized.




Fig. 27—PBT store test set and digital strobe set connected to a store.

X. PHYSICAL DESIGN

10.1 *Frame Description*

The store physical design uses No. 1 ESS style frameworks and packaging techniques⁷ wherever possible. The store physical design problem can be stated by four general requirements:

- (i) Package the store circuitry within a single 7-foot high double-bay framework.
- (ii) Separate circuit functions to eliminate noise interference that would cause erroneous circuit operations.
- (iii) Locate circuit functions to minimize interconnecting lead lengths.
- (iv) Provide adequate thermal dissipation for high wattage components and isolate them from temperature sensitive circuitry.

The store, shown in Fig. 1, consists of one double-bay framework that contains four piggyback twistor (PBT) modules; access, read, and write circuits; and other related circuitry. The location of the major circuits is shown in Fig. 28. ESS No. 1 type circuit packs are used to package the majority of the store circuitry. Over 600 of these boards, representing 58 different codes, are required.

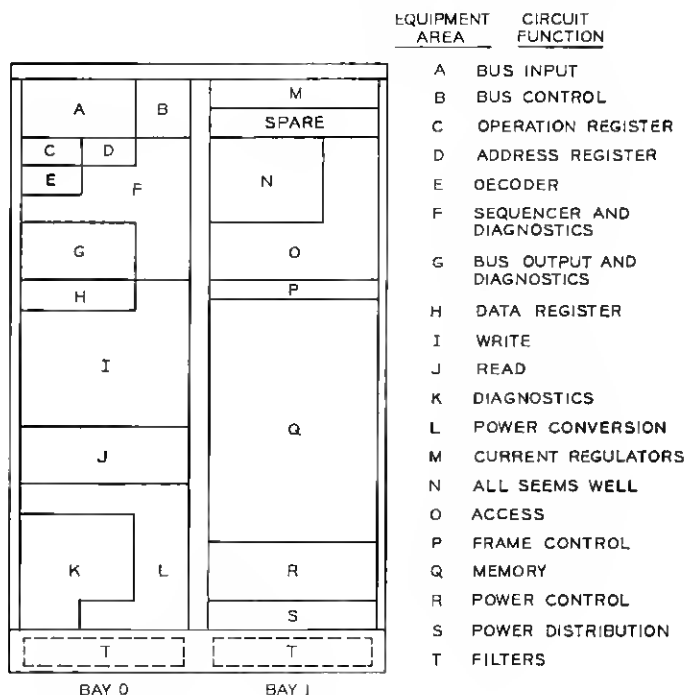


Fig. 28—SPC No. 1A store showing the location of the major circuits.

The requirement to package the store circuits within a single double-bay framework requires that the PBT modules be mounted within 34 inches of bay space. This objective is accomplished with a back-to-back mounting of the four modules, which also shortens interconnecting cable lengths. Of particular importance was the need to separate the access circuit, which produces 2.0 ampere pulses, from the logic, register and read circuitry. To this end, bay 1 contains the high level access and power distribution circuits along with the four back-to-back PBT modules. All other store functions are contained within bay 0.

Five current regulators that supply access and bias current for the modules contain components operating at temperatures up to 250°F. The regulators, in accordance with requirement (iv), are mounted at the top of bay 1 on a single mounting plate. A 4-inch wide unblocked area is retained below the regulators to allow convective cooling of

regulator components. A detailed description of these regulators is given in Section 10.3.1.

10.2 *Power Dissipation*

Each SPC No. 1A store frame in an idle or read state dissipates 1240 watts. An additional 585 watts may be introduced by writing at the rate of one write operation every 157 μ s in a particular store. This represents the highest average writing frequency of present SPC No. 1A applications. Presently a maximum of 40 stores, dependent on memory requirements, may be provided with the SPC No. 1A. Thus, the store has considerable effect on office ventilation⁸ and power plant requirements.

Approximately 62 percent of the total power dissipation is concentrated in three areas; the five current regulators, the four PBT modules and the write circuit packs. Each of two read current regulators dissipates approximately 75 watts, each of two write regulators 95 watts, and the bias regulator 135 watts. For the highest expected write rate, each of the four PBT modules dissipates 40 watts, including 20 watts of bias. The write circuit consumes 505 watts. Power levels of these magnitudes require special design treatment to avoid high component operating temperatures which increase failure rates. The following section deals with the thermal design considerations in some detail.

10.3 *Thermal Characteristics*

The thermal characterization of the store was achieved by four laboratory experiments involving the 15B modules, the current regulators, and the circuit packs within bay 0 of the frame.

10.3.1 *Current Regulator*

Five current regulators are required to supply bias and read and write access current for the PBT modules. The regulator circuit is an active series type regulator utilizing a power transistor and resistors as the series elements. A plug-in unit design was adopted because it offers the advantage of a fully tested and adjusted functional unit. Also, it is readily accessible for maintenance or replacement in a working store with minimal downtime. Of primary importance for a final overall unit design is the high power dissipation that is a characteristic of a series type regulator circuit. The three design objectives needed to accommodate this high power dissipation are:

(i) Provide efficient heat-transfer mechanisms to limit component operating temperatures to safe values.

(ii) Isolate power dissipating components from thermally sensitive regulator components.

(iii) Limit the exposed surface temperature of the package to 140°F to prevent personnel injury when extracting the regulator.

A laboratory evaluation of models indicated an economical approach to produce a plug-in unit (see Fig. 29) that meets these three objectives. A conventional heat sink was found adequate to limit the transistor temperature within its rating. One quarter inch thick black lacquered aluminum parts were used to construct a rectangular heat sink mounting for the resistors. To achieve requirements (ii) and (iii) a package design utilizing six thermally isolated areas within the regulator was conceived and evaluated. The principle form of heat transfer between the six areas is conduction, whereas the principle form of heat transfer from the unit is convection.

Laboratory evaluation of regulator models of the compartmented design indicated that only semi-isolation between the six areas was essential. Therefore, only some degree of restriction of conductivity was




Fig. 29—Current regulator.

necessary to accomplish the objectives. This thermal restriction was accomplished by reducing the total area of contact between sections to four small connecting areas. This reduced cross-sectional area between sections established a thermal gradient large enough to greatly reduce the thermal effects of the resistors and transistor on each other and on the more sensitive components in the regulator. Perforated steel, 50 percent open, further increased the temperature gradient along the right side plate to the front panel.

Power resistors and transistors, which normally operate at elevated surface temperatures, were located internally to minimize danger to personnel. Heat from these components is dissipated by conduction to heat sinks to insure that temperature limits are not exceeded. Somewhat in opposition to this arrangement is the requirement that outside surfaces of the unit must not be a hazard to maintenance personnel. Accordingly, the design provides thermal isolation for the front panel and right side plate which limits these surfaces to 135°F even under maximum operating conditions. This is within the 140°F design limit specified by objective (iii).

10.3.2 15B PBT Module

The 15B module (see Fig. 2) contains two heat sources, the digit lines, and the bias winding. The principle form of heat transfer from the module is by convection from the front surface of the memory planes and from vertical fins extending outward from the rear of the module. Each fin covers the area of a memory plane and extends past the edge of the plane by 0.65 inch. Thus a 0.65 by 12 inch area of each fin extends into a chimney formed by the back-to-back mounting of the modules.

Since the maximum internal temperature of the module must not exceed 180°F, a study was conducted to establish the relationship between module power dissipation and its internal temperature. Write rate was used as the variable parameter instead of power to provide an easily understood expression that can be directly applied by system programmers who determine store usage.

The write rate (t_p) is defined as the percent of the time required to perform a write operation in a particular store (50.4 μ s) with respect to the average time between the start of consecutive write operations (τ). The definition may be represented by the following formula:

$$t_p = \frac{50.4 \times 10^{-4}}{\tau}.$$

Module thermal properties were established through a series of tests

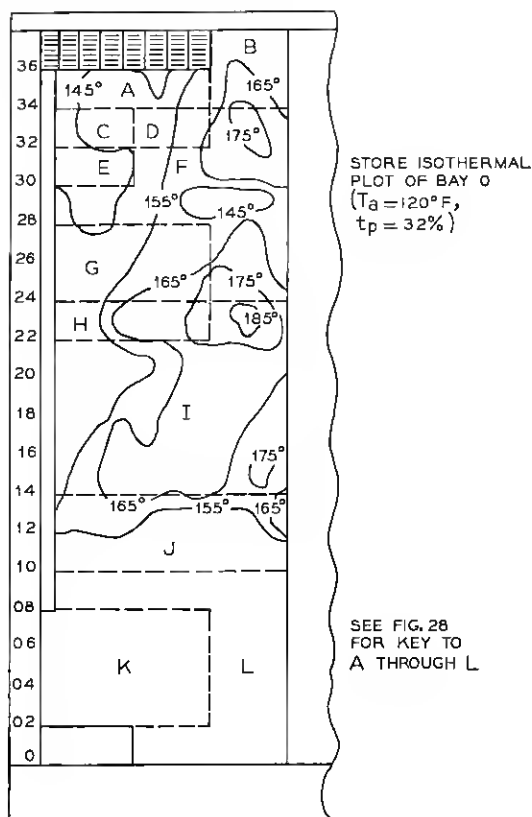


Fig. 30—Frame isotherm plot.

on a store frame maintained at a constant ambient temperature in an environmental chamber. These tests established the relationship between write rate and module hot spot temperature as:

$$T_{\max} = 0.78t_p + T_a + 13.$$

where T_a is the ambient temperature.

The internal temperature limit of 180°F restricts continuous store operation to a 60.2 percent write rate at the ambient temperature limit of 120°F . Knowledge of this write rate limitation is an important application parameter. It represents a relationship between the method of use and the component physical characteristics which always exists but is often unknown. Considerable temperature safety margins exist

since present SPC No. 1A applications will not exceed a 32 percent write rate.

10.3.3 Bay 0 Circuit Packs

The circuit packs mounted within bay 0 of the store frame dissipate approximately 770 watts. To determine component operating temperatures the store was placed in an environment whose temperature was near the maximum ambient and with air movement by natural convection only. The store was then exercised at the expected worst case write rate of 32 percent ($\tau = 157\mu\text{s}$). After the store reached thermal equilibrium, temperatures were measured to produce an isothermal plot. This plot, corrected to the maximum ambient of 120°F , is shown in Fig. 30.

The resulting maximum component environmental temperature was found to be 185°F for an ambient temperature of 120°F . No component was found to operate above its maximum temperature ratings. Since the thermal characterization was produced without forced air movement, the air velocity about the frame was due only to natural convection and represents a worse case condition. With some forced air motion, as now recommended for these equipments, additional temperature margins are provided.

XI. SUMMARY

The requirements of the SPC 1A store motivated the development of a new and highly sophisticated memory device. The complexity of this device, the PBT memory, led to an intensive test program in order to understand the theory and to characterize its performance. The results of this effort have produced a memory that is very well characterized and stringently tested.

The resulting store combines the nondestructive readout features of the No. 1 ESS program store and the electronic alterability of the call store. The preread write sequence along with the protected area feature has made accidental program destruction practically impossible. The electronic alterability makes office data and program changes an easy, quick procedure. Field experience has demonstrated data integrity over wide ranges of store adjustments.

REFERENCES

1. James, D. A., Myers, F. H., "An Electrically Alterable NDRO Memory System. Using a Twistor," IEEE Trans. Communication Technology, COM-17 (April 1969), pp. 93-99.

2. Gerke, P. M., Harding, P. A., and Staehler, R. E., "No. 1 ESS Call Store—A 0.2-Megabit Ferrite Sheet Memory," B.S.T.J., 43 No. 5 (September 1964), pp. 2147-2191.
3. Ault, C. F., Gallaher, L. E., Greenwood, T. S., and Koehler, D. C., "No. 1 ESS Program Store," B.S.T.J. 43 No. 5 (September 1964), pp. 2097-2146.
4. Baker, W. A., "The Piggyback Twistor—An Electrically Alterable Nondestructive Readout Twistor Memory," IEEE Trans. on Comm. and Elec., 83 (November 1964), pp. 829-833.
5. Nesbitt, E. A., Chin, G. Y., and Jaffe, D., "New Low-Magnetostrictive Permanent-Magnet Alloys," J. Appl. Phys., 39 (February 1968), pp. 1268-1269.
6. Myers, F. H., "Using a Small Computer to Test a Complex Memory System," Proceedings of the National Electronics Conf., Chicago, Illinois, 25 (December 1969), pp. 711-715.
7. Ferguson, J. G., Grutzner, W. E., Koehler, D. C., Skinner, R. S., Skubiak, M. T., and Wetherell, D. H., "No. 1 ESS Apparatus and Equipment," B.S.T.J., 43 No. 5 (September 1964), p. 2355.
8. Baldinger, W. R., and Clark, G. T., "TSPS No. 1: Physical Design," B.S.T.J., this issue, pp. 2685-2709.